



# Lustre 2.0 Alpha-7 Summary

## Summary

The focus of the seventh Alpha release (Alpha-7) of 2.0 was to continue to improve stability of Lustre through landing of fixes to the main trunk of the Lustre code base. 157 total fixes were landed during this cycle.

This release was tested on RHEL5/x86\_64 and OEL5/x86\_64 with both IB and TCP connectivity. Known failures are documented in the HEAD Daily Testing Document available from [lustre.org](http://lustre.org).

## Timeline

The Alpha-7 bug fix period was from November 25, 2009 – February 5, 2010. The build on February 5, 2010 marks the Lustre 2.0 Alpha-7 release which was announced and available for download on February 10, 2010.

## Fixes Landed

Below is a summary listing of all unique landings to HEAD during the Lustre 2.0 Alpha-7 Milestone period. This is inclusive of all landings -- code fixes as well as test framework fixes.

1. Bug 15332: Lack of congestion control in LNET causing timeouts in routed configuration. There are 2 references to landings but no patch attached in the appropriate date range.
2. Bug 15692: replay-dual test\_18 Watchdog triggered  
There are 2 references to landings but no patch attached to the bug.
3. Bug 15864: insanity test\_6 (FAILURE\_MODE=HARD): The ost\_connect operation failed with -16  
Attachment: 27754
4. Bug 15977: conf-sanity test\_39: @@@@ @@@@ FAIL: memory leak detected  
Attachment: 27959 - Landing flag is not set but there is a record of the landing in the logs.
5. Bug 16186: One down Lustre FS hangs ALL mounted Lustre filesystems  
Attachment: 25054
6. Bug 16466: review quota recovery  
Attachment: 27341
7. Bug 16715: NFS cannot open files in HEAD, MDS crashing  
Attachment: 27697
8. Bug 16721: 4\*racer.sh: osd\_object\_ref\_del() ASSERTION(inode->i\_nlink > 0)  
Attachment: 27244

9. Bug 16890: small fixtures for quota  
Attachment: 26551 Attachment: 26605
10. Bug 17167: libcfs: ensure all libcfs exported symbols to have cfs\_ prefix  
Attachment: 28001
11. Bug 17545: open sometimes returns ENOENT instead of EACCES  
Attachment: 27392
12. Bug 17682: Improve debug capability for production systems  
Attachment: 27496 - There are 2 landing references in the logs but only 1 landed patch in the bug report.
13. Bug 18357: Down OSTs block all MDT threads resulting in client hangs  
Attachment: 26410
14. Bug 18460: Asymmetric Failure of a Router  
There's a reference in the logs but no patches landed in the bug.
15. Bug 18630: LustreError: 9153:0:(quota\_context.c:622:dqacq\_completion()) LBUG  
Attachment: 27701
16. Bug 18690: smaller hash bucket sizes, cleanups  
Attachment: 28140
17. Bug 18790: Build failure on HEAD, RHEL 5, PPC64 Client  
Attachment: 27495
18. Bug 19023: client isn't evicted as expected by test  
Attachment: 27269
19. Bug 19325: ldlm\_extent\_internal\_policy\_waiting() causing soft lockups  
Attachment: 26003 - Patch later reverted.
20. Bug 19387: integrate LST into acc-sm  
Attachment: 28168
21. Bug 19518: Move lustre hashes to libcfs  
Attachment: 27126
22. Bug 19556: replay-single test\_53c: @@@@ @ @ FAIL: test\_53c failed with 2  
Reference to a landing in the logs but no patch attached.
23. Bug 19557: BUG: soft lockup - CPU#1 stuck for 10s! [ll\_mdt\_07:4523]  
Attachment: 27357 Attachment: 27408
24. Bug 19616: t-f except list file  
Attachment: 27123
25. Bug 19669: Store HSM flags on MDT and release  
Attachment: 27308
26. Bug 19673: Add ldiskfs maxdirsize mount option  
Attachment: 27654
27. Bug 19689: mkfs.lustre, tuneufs.lustre --mountfsoptions handing is poorly designed  
Attachment: 23782
28. Bug 19702: runttests : @@@@ @ @ FAIL: old and new files are different 11  
Attachment: 27779
29. Bug 19720: various -Werror illicited errors  
Attachment: 27957
30. Bug 19721: debian/ubuntu .deb build bits  
Attachment: 27138
31. Bug 19735: multiple-instance support for kptllnd  
There is reference in the logs to a landing but there is no patch attached to the bug.

32. Bug 19742: sanity test\_130a/b/c/e: @@@@ @ FAIL: FIEMAP on 1-stripe file(/mnt/lustre/f130a) failed;  
Attachment: 27769 Attachment: 27719
33. Bug 19808: 2.6.29-fc11 patchless client support  
Attachment: 26955 Attachment: 27780 - Later reverted. Attachment: 27853
34. Bug 19813: ASSERTION(atomic\_read(&client\_stat->nid\_exp\_ref\_count) == 0) failed:count 61  
Attachment: 27946
35. Bug 19908: lustre get xattr bug  
Attachment: 27240
36. Bug 19964: SOM EA  
Attachment: 24582 Attachment: 24593 Attachment: 26265 Attachment: 26266
37. Bug 20057: Autovetting and test-framework enhancements  
Attachment: 27181
38. Bug 20098: mod use after free  
Attachment: 27434
39. Bug 20101: sanity.sh 27w does not work as expected, but still does not fail  
Attachment: 27197
40. Bug 20171: clients failed to detect downed routers  
Landing referenced in the logs but no patch attached.
41. Bug 20237: acc-sm: recovery-\*-scale tests gather logs fn  
Attachment: 26515
42. Bug 20301: unable to perform "mount -t lustre" of 16TB OST device  
Attachment: 27791
43. Bug 20315: Use libexecdir for scripts  
2 Landings referenced in the logs but no patch attached.
44. Bug 20580: OOM failure in sanity, test 78  
Attachment: 26844
45. Bug 20583: metabench hang due to a dump on MDS  
Attachment: 27582
46. Bug 20591: race on obd\_device::md\_stats  
Attachment: 27266 Attachment: 27762
47. Bug 20607: hyperion: Oops during metabench  
Attachment: 27615
48. Bug 20617: Add mptlinux and nxge drivers to Lustre builds  
Attachment: 27208 Attachment: 27206 Attachment: 27329 Attachment: 27365 - later reverted.
49. Bug 20645: LBUG: (lov\_io.c:256:lov\_page\_subio()) ASSERTION(lio->lis\_cl.cis\_io != NULL) failed  
Attachment: 27396 - Later reverted. Attachment: 27710
50. Bug 20655: enable ext4 by default for 2.0  
Attachment: 27202 Attachment: 27348
51. Bug 20722: Fix watchdog timer message to be more clear  
Attachment: 27497
52. Bug 20758: Kernel update for SLES10 SP2 2.6.16.60-0.42.8  
Attachment: 27918
53. Bug 20816: replay-single test\_67b: @@@@ @ FAIL: cp failed : EIO  
Attachment: 26749

54. Bug 20853: request portal in lustre\_idl.h for Cray DVS  
Attachment: 26956
55. Bug 20863: OST pools testing does not clean up after itself  
Attachment: 26686
56. Bug 20866: lustre\_rmmod: DEPS assignment needs quotes  
Attachment: 26183
57. Bug 20892: bad entry in directory xxx: inode out of bounds  
Attachment: 27190
58. Bug 20918: recovery-mds-scale: Failed to meet interval 5 times ( REQFAIL=4 ); have sleep=-71  
Attachment: 27925
59. Bug 20953: sanity-quota test\_30: @@@@ @@@@ FAIL: grace times were reset or unexpectedly high latency  
Attachment: 26296
60. Bug 20970: write\_disjoint failed  
Attachment: 27941
61. Bug 20989: (lov\_merge.c:74:lov\_merge\_lvb()) ASSERTION(spin\_is\_locked(&ism->ism\_lock)) failed (SR 71691004)  
Attachment: 28212
62. Bug 21041: sg\_set\_page autoconf macro needs asm/types.h  
Attachment: 26434
63. Bug 21079: sanity.sh: FAIL: test\_215 /proc/sys/net/buffers does not look as expected  
Attachment: 27376
64. Bug 21103: ORNL LCE Router features\fixes  
Attachment: 26666 Attachment: 26694 Attachment: 26716 Attachment: 26757
65. Bug 21115: test didn't get run but yala report it as failed  
Attachment: 26736
66. Bug 21128: racer: BUG: unable to handle kernel paging request  
Attachment: 27258
67. Bug 21142: cascading\_rw has to be skipped for NFSCLIENT=yes  
Attachment: 26668 Attachment: 27432
68. Bug 21156: dead code in socklnd\_lib-linux.c  
Attachment: 26945
69. Bug 21224: running sanity\_quota or OST pools test 23 from acc-sm on b1\_8 causes reboot of test system  
Attachment: 27287
70. Bug 21240: replay-single test\_44a: error: recover: Connection timed out  
Attachment: 27262
71. Bug 21259: lfs check is root-only in 1.8  
Attachment: 27035 Attachment: 27906
72. Bug 21264: recovery-mds-scale: Bus error dd bs=4k count=1000000 if=/dev/zero of=\$TESTDIR/dd-file  
Attachment: 28245
73. Bug 21265: recovery-mds-scale need bigger timeout  
Attachment: 27088
74. Bug 21275: rename lreplicate to lustre\_rsync  
Attachment: 27140

75. Bug 21314: genops.c:1019: \_\_class\_export\_del\_lock\_ref()) ASSERTION(lock->l\_exp\_refs\_nr > 0) failed  
Attachment: 27010
76. Bug 21348: Timer expired but recovery does not end (b1\_8)  
Attachment: 27378
77. Bug 21355: ASSERTION(cfs\_fs\_pwd(current->fs) == new\_ctx->pwd) failed  
Attachment: 28288
78. Bug 21370: sanity test\_27x fail: OST0 was degraded but new created file still use it  
Attachment: 28109
79. Bug 21372: sanity-quota test\_28 hang  
Attachment: 27285
80. Bug 21375: lustre/obdclass/obd\_mount.c:2104: error: 'FS\_HAS\_FIEMAP' undeclared here (not in a function)  
Attachment: 27261
81. Bug 21426: apply test patches to OFED during lbuild  
Attachment: 27204
82. Bug 21430: requests handled before \*\_CONNECT  
Attachment: 27270
83. Bug 21433: sanity-quota test\_3: @@@@ @@@@ FAIL: write success, but expect EDQUOT  
Attachment: 27753 Attachment: 28118
84. Bug 21434: sanityN test\_32b: @@@@ @@@@ FAIL: lockless i/o was not triggered  
Attachment: 27237
85. Bug 21448: mdsrate-create-large test\_4: @@@@ @@@@ FAIL: mdsrate creates for a single client failed, aborting  
Attachment: 27718
86. Bug 21454: cl\_io\_prepare\_write()) ASSERTION(cl\_page\_in\_io(page, io)) failed  
Attachment: 27529
87. Bug 21457: download code in lbuild refactoring  
Attachment: 27305
88. Bug 21459: Peer state incorrect on routers  
Attachment: 27306
89. Bug 21460: ll\_commit\_write() ASSERTION(result != -EALREADY) failed  
Attachment: 27253
90. Bug 21489: cp -p command does not preserve the dates and timestamp  
Attachment: 28006 - Later reverted. Attachment: 28133
91. Bug 21490: 1.8<->2.0 interop: in 2.0 server and 1.8 client env. read 1.8 file hung  
Attachment: 28095
92. Bug 21499: LBUG: (mdt\_reint.c:111:mdt\_version\_get\_check()) ASSERTION(mdt\_object\_exists(info->mti\_mos[index])) failed  
Attachment: 27983
93. Bug 21513: Client reconnects only to failover NID never primary  
Attachment: 27334
94. Bug 21514: lbuild should not call the ccache -s exit trap if interrupted  
Attachment: 27337
95. Bug 21516: lreplicate-test.sh script is missing  
Attachment: 27386

96. Bug 21519: conf-sanity test 16, OST general protection fault in lprocfs\_counter\_add  
Attachment: 27794
97. Bug 21527: use RPMs %configure macro in our lustre.spec file  
Attachment: 27448 - Later reverted.
98. Bug 21530: SLES10 ppc64 patchless client failes to build  
Attachment: 27430
99. Bug 21543: e2fsck: MMP: invalid magic number while trying to open /dev/lustre/ost4  
Attachment: 27702 Attachment: 27714
100. Bug 21544: t-f does not init gss if started on already mounted lustre  
Attachment: 27641
101. Bug 21547: add cascading\_rw.c to lustre/tests  
Attachment: 27650 Attachment: 27737
102. Bug 21551: racer, (cl\_page.c:233:cl\_page\_gang\_lookup()) slice != NULL  
Attachment: 27752
103. Bug 21558: replay-dual test\_21b: @@@@ @@@@ FAIL: Not all renames are replayed.  
COS=0  
Attachment: 27772
104. Bug 21564: e2fsck should warn when MMP update interval is extended  
Attachment: 28105
105. Bug 21565: filter\_last\_id() NULL deref  
Attachment: 27985
106. Bug 21571: Loadgen improvements  
Attachment: 27456 Attachment: 27559
107. Bug 21586: clarify lbuild error reporting  
Attachment: 27493 Attachment: 27494 Attachment: 27657
108. Bug 21593: fix sanityN test\_33 to find correctly proc jbd file for ext[3|4]  
Attachment: 27510
109. Bug 21595: mdsrate-create-large.sh, BUG: soft lockup - CPU#0 stuck for 10s!  
[ll\_ost\_io\_04:8153]  
Attachment: 28047 Attachment: 28087
110. Bug 21600: SLES11 i686 & x86\_64 patchless fail to build on HEAD  
Attachment: 27516 Attachment: 27708
111. Bug 21632: Kernel update for RHEL5.4 2.6.18-164.11.1.el5.  
Attachment: 28122
112. Bug 21656: Quiet CERROR("dirty %d > system dirty\_max %d\n")  
Attachment: 27820
113. Bug 21670: 1.8.2 RHEL5 & SLES10 ppc64 patchless builds fail  
Attachment: 27698 Attachment: 27739 Attachment: 27748  
Attachment: 27749 - Later reverted. Attachment: 27750 Attachment: 27757
114. Bug 21671: RHEL5/ext4 builds fail on patched  
Attachment: 27706
115. Bug 21709: recovery-small test\_23 failed, LBUG, (mdd\_device.c:1287:md\_quota())  
ASSERTION(env->le\_ses != NULL) failed  
Attachment: 27786
116. Bug 21716: sanity test\_78: @@@@ @@@@ FAIL: rdwr failed  
Attachment: 27915
117. Bug 21746: The client profile could not be read from the MGS  
Attachment: 27860

118. Bug 21759: miscellaneous lbuild cleanups  
Attachment: 27855
119. Bug 21828: OSS stops handling requests during recovery  
Attachment: 28125
120. Bug 21846: LBUG on null key while rehashing quota lqs  
Attachment: 28161
121. Bug 21879: MDS Oopses in ptlrpc\_dispatch\_difficult\_reply()/\_\_list\_add()  
Attachment: 28246
122. Bug 21882: lru resize SLV can get stuck  
Attachment: 28138
123. Bug 21889: Can't find test script for lustre-rsync-test 2 landings are recorded in the logs but there are no patches attached to the bug.
124. Bug 21890: replay\_dual: mds umount failed  
A landing was referenced in the logs but no patch is attached to the bug.
125. Bug 21924: lustre-rsync-test test\_7: @@@@ @@@@ FAIL: Failure in replication; differences found.  
Attachment: 28291

## Results

Please review the Lustre 2.0 Alpha Test Plan document and Head Daily Testing Results spreadsheet (both available on [lustre.org](http://lustre.org)) for detailed plans and daily testing results.

The Lustre 2.0 Alpha-7 exhibits the following characteristics:

- Testing completed for RHEL5/x86\_64 and OEL5/x86\_64 using IB and TCP
- Passing consistently: all daily tests are passing consistently on small clusters
- Known Issues: NFS (bugs 22040, 22042, 21678); interoperability between Lustre 1.8.x and 2.0 does not work with this Alpha version

Lustre 2.0 Alpha-7 was delivered on February 10, 2010 and is available from:

<http://downloads.lustre.org/public/lustre/v2.0/latest/>

or

[http://downloads.lustre.org/public/lustre/v2.0/alpha/Lustre\\_2.0\\_Alpha7/](http://downloads.lustre.org/public/lustre/v2.0/alpha/Lustre_2.0_Alpha7/)